arXiv:1406.5062v1 [stat.CO] 19 Jun 2014

Overcoming computational inability to predict clinical outcome from high-dimensional patient data using Bayesian methods

A. SHALABI A. C. C. COOLEN

Institute for Mathematical and Molecular Biomedicine
King's College London
London, U.K
akram.shalabi@kcl.ac.uk
ton.coolen@kcl.ac.uk

E. de RINALDIS

NIHR Biomedical Research Centre
R&D Department
Guy's Hospital
London, U.K.
emanuele.de_rinaldis@kcl.ac.uk

Abstract

Clinical outcome prediction from high-dimensional data is problematic in the common setting where there is only a relatively small number of samples. The imbalance causes data overfitting, and outcome prediction becomes computationally expensive or even impossible. We propose a Bayesian outcome prediction method that can be applied to data of arbitrary dimension d, from 2 outcome classes, and reduces overfitting without any approximations at parameter level. This is achieved by avoiding numerical integration or approximation, and solving the Bayesian integrals *analytically*. We thereby reduce the dimension of numerical integrals from $2d$ dimensions to 4, for any d. For large d, this is reduced further to 3, and we obtain a simple outcome prediction formula without integrals in leading order for very large d. We compare our method to the *mclustDA* method (Fraley and Raftery 2002), using simulated and real data sets. Our method perform as well as or better than *mclustDA* in low dimensions d. In large dimensions d, *mclustDA* breaks down due to computational limitations, while our method provides a feasible and computationally efficient alternative.

Keywords: Discriminant analysis; Bayesian outcome prediction; Overfitting; Curse of dimensionality; Bayesian integration in high dimensions; Binary-class prediction.

1 Introduction

Outcome prediction is based on discriminant analysis, which is the use of known classifications of some observations to classify others [1]. The aim is to use some observations in the data set under study as the training set, to automatically discover and learn the rules linking an observation to its class. These rules can then be applied to the remaining observations in the test or validation set to assign them to classes. Discriminant analysis methods are often probabilistic, or model-based, where the observations in each class are assumed to be generated by a distribution specific to that class [1].

Fraley and Raftery [2] developed a popular model-based discriminative analysis approach to outcome prediction that has been used in a variety of contexts [3, 4, 5, 6, 7]. This involves producing Bayesian posterior

probabilities of belonging to a class, or class-conditional probabilities, by using model-based clustering to fit a Gaussian mixture model (GMM) to each class in the training set, and estimating the optimal parameters whilst ignoring the uncertainty in these estimates. These are then used to calculate the class-conditional probabilities of the remaining observations.

Using leave-one-out cross-validation (LOOCV) [8, 9, 10], we found that whilst their approach produces accurate and efficient results for data sets consisting of a limited number of samples with low dimensionality, a number of challenges arise for data sets with large dimensions. These are data overfitting [11, 12], and computational inability [11]. These factors combined produce outcome predictions of limited reliability if any at all. Each parameter contributes a d-dimensional Bayesian integral to the predictive distribution, which is at best computationally expensive to estimate numerically. Instead, at the expense of ignoring the uncertainty in the parameters, they are estimated by the most probable values using the Expectation-Maximisation algorithm (EM) [13, 2]. Yet this is also computationally expensive in high-dimensions [2, 14, 15]. Several dimension-reduction approaches such as principal component analysis (PCA) [16, 2, 17], subspace clustering [18, 14], and/or using constrained and parsimonious models [15], have been suggested in literature. The first approach runs the risk of compromising the prediction accuracy through loss of information found in higher dimensions [19, 20]. The second assumes that high-dimensional data live in subspaces with dimensionality less than d [18, 14, 21] but relies on finding a good classifier in high dimensions [21, 20], and the third approach is a trade-off between perfect modelling and what can be correctly estimated in practice. The latter two have been combined to reduce the number of parameters and dimensions for EM to consider [14, 22, 20]. However underestimated parameter uncertainty remains.

The solution to these challenges will involve some form of dimension reduction without underestimating parameter uncertainty or loss of information in the data. In this paper we show that it is possible to do so by solving the d-dimensional Bayesian integrals analytically, and propose a Bayesian outcome prediction method that can be applied to any arbitrary number of dimensions d without any approximations at parameter level. In Section 2, we briefly describe the method developed by Fraley and Raftery [2], and in Section 3 we describe our Bayesian outcome prediction method. In Section 4, we apply our method to simulated examples, and show how the curse of dimensionality [23] is lifted in terms of CPU and memory. In Section 5, we apply our method to a real biomedical example. We conclude with a discussion in Section 6.

2 Model-based discriminant analysis

Fraley and Raftery [2] developed a model-based approach to discriminant analysis which is implemented in the R software environment [24] using the reference package *mclust* [2, 17]. In this section, we briefly describe their method.

Many methods have been proposed for discriminant analysis, and it is applicable in a wide variety of settings [25, 26, 27]. Discriminant analysis methods are often probabilistic, based on the assumption that observations in each class are generated by a distribution specific to that class. The number of classes J is assumed to be known. We typically have a multi-dimensional data set D consisting of N observations, and their corresponding class labels $\sigma \in \{0, \ldots, (J-1)\}$ that indicate which class each observation belongs to: i.e. $D = \{(\boldsymbol{x}_1, \sigma_1), \ldots, (\boldsymbol{x}_N, \sigma_N)\}$, consisting for each individual i, of a d-dimensional covariate vector $\boldsymbol{x}_i$, and the class σ_i to which that individual belongs to. The data set is then split into a training and test set. Each pair $(\boldsymbol{x}_i, \sigma_i)$ is assumed to be is drawn independently, from an unknown joint probability distribution $p(\boldsymbol{x}, \sigma)$, that describes the data set being studied. In biomedical data sets, we typically have $d \gg N$ [11, 12, 28, 29, 30, 31, 32]. If we take $p(\boldsymbol{x}_i|\sigma)$ to be the probability distribution of the observations belonging to class σ, and $p(\sigma)$ as the proportion of observations that belong to that class, or class-imbalance parameter, then according to Bayes' theorem, the posterior probability that an observation $\boldsymbol{x}_i$ belongs to class σ, or class-conditional probability, is

$$p(\sigma|\boldsymbol{x}_i) = \frac{p(\boldsymbol{x}_i|\sigma)\ p(\sigma)}{\sum_{\sigma=0}^{J-1} p(\boldsymbol{x}_i|\sigma)\ p(\sigma)} \tag{1}$$

In probabilistic discriminant analysis, a model is fit to each class in the training set, and observations in the test set are assigned to the class corresponding to the model in which they have the highest class-conditional probability. Most common model-based discriminant analysis methods are based on the assumption that the observations in each class are multivariate Gaussian [33, 34]. GMMs are also commonly used [1, 2, 15, 20].

Fraley and Raftery [2] fit one GMM to each class in the training set using model-based clustering, and applied the optimal model to the test set. This withdrew the need to know the number of classes, and covariance matrix structure a priori by allowing the data in the training set to determine the number of classes in the data set. The procedure is called *mclustDA* [2, 17], and can be summarised as follows: Set the number of GMM components, and type(s) of covariance matrix structure to consider. Apply model-based hierarchical clustering [2, 35, 33, 34] to the training set to provide initial, suboptimal partitions. These provide initial values for the parameters and class-conditional probabilities, which are used to initialise the Expectation-Maximisation algorithm (EM) [13, 34, 33]. EM is applied to find the optimal parameters for each GMM, assigning a class-conditional probability to each data point. The Bayesian Information Criterion (BIC) [36, 33, 34, 2] is then computed for each GMM, with the one with the highest BIC estimated to be the optimal model. Degenerate solutions are alleviated by using the posterior mode (MAP) instead of the maximum likelihood estimates (MLE) for the EM and BIC calculations [4]. The optimal model is then applied to the test set, and the class-conditional probabilities are calculated using (1). An observation in the test set is assigned to the class corresponding to the highest class-conditional probability. It has been suggested that *mclustDA* may not work well or be insufficient for high-dimensional data due to computational inability and/or overfitting [2, 37, 38, 15, 17]. A suggested solution is to restrict the parameters of the model [15, 17] or apply a dimension-reduction technique such as PCA [17]. We apply the former at the risk of fitting an inadequate model to the data rather than risk losing information in the data by using an a priori dimension reduction technique.

We fit one isotropic Gaussian distribution to each class such that $\boldsymbol{x}|\sigma \sim \mathcal{N}(\boldsymbol{\mu}_\sigma, \lambda_\sigma^2 \mathbb{I}_{d\times d})$ where $\boldsymbol{\mu}_\sigma$ is a $d \times 1$ mean vector, and λ_σ^2 is the variance of the distribution for class σ. This is achieved by choosing **EII** (Equal (**E**) volume, Identity Matrix (**I**) shape and Identity Matrix (**I**) orientation) as the covariance structure, *mclustDA* as the model type, *defaultPrior* as the prior, and setting G (the number of Gaussians to be considered per class) to 1 [2, 17]. Using LOOCV, we found that for data sets consisting of a limited number of samples ($N = 100$) with $d = 3, 10, 100, 1000$, *mclustDA* produces predictions in a reasonably quick time with little computational expense. For data sets consisting of a large number of samples with low d, *mclustDA* was computationally expensive due to the large amount of time taken by model-based hierarchical clustering [2, 37, 38]. Taking a random subset of the data set as the training set, and classifying the remaining observations (in reasonably sized blocks) [35, 34, 2, 17] is more efficient, but can lead to decreasing accuracy [10, 37, 38]. A more accurate incremental approach has been developed but only for data sets whose storage and memory requirements are not too large to be handled in whole by the *mclust* software package [37, 38]. For large d (i.e. $d = 10,000$), *mclustDA* could not handle the large dimensionality of the data set, and failed to produce any results. We surmise that for a large number of samples, it would be extremely expensive in terms of CPU power and memory to produce any results given the combined computational expense of model-based hierarchical clustering, and EM. For a limited number of samples, we would be overfitting heavily. We describe a method that both overcomes this computational inability, and reduce overfitting without sacrificing any information contained in the data set in Section 3.

3 Bayesian binary outcome prediction for any number dimensions

In this section, we present a Bayesian outcome prediction method that can be applied to 2 homogeneous classes, and any number of dimensions. We reduce the dimension of numerical integrals from $2d$ dimensions to 4, for any d (FD). For large d, this is reduced further to 3 (LD), and we obtain a simple outcome prediction formula without integrals in leading order for very large d (LOLD).

We fit one isotropic Gaussian distribution to each class such that $\boldsymbol{x}|\sigma \sim \mathcal{N}(\boldsymbol{\mu}_\sigma, \lambda_\sigma^2 \mathbb{I}_{d\times d})$ where $\boldsymbol{\mu}_\sigma$ is a $d \times 1$ mean vector, and λ_σ^2 is the variance of the distribution for class σ. The parameter prior with respect to the mean is set to be $\boldsymbol{\mu}_\sigma \sim \mathcal{N}(0, \alpha_\sigma^{-1} \mathbb{I}_{d\times d})$, where α_σ is the precision of the distribution, and use $\delta(\lambda_\sigma - \psi_\sigma)$ to fix λ_σ, and treat it as hyperparameter. We only take the uncertainty at parameter level into account. If we know precisely or wish to control the class-imbalance in the data set, we can fix $p(\sigma)$ to our chosen value σ'_σ. In the absence of any information about the class-imbalance in the data set, we take this to be the class-imbalance in the training set $\bar{\sigma}_\sigma = N_\sigma/N$ (the maximum entropy distribution for $p(\sigma)$), where N_σ is the number of observations in class σ in the training set [21]. Proceeding with $\sigma = 1$ ($p(0|\boldsymbol{x}^\star, D)$ can be calculated by subtracting $p(1|\boldsymbol{x}^\star, D)$ from 1), the predictive distribution for a new observation $\boldsymbol{x}^\star$ is

$$p(1|\boldsymbol{x}^\star, D) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathrm{d}\boldsymbol{\mu}_0 \mathrm{d}\boldsymbol{\mu}_1\, p(\boldsymbol{\mu}_0|D, \hat{\psi}_0, \hat{\alpha}_0)\, p(\boldsymbol{\mu}_1|D, \hat{\psi}_1, \hat{\alpha}_1)\, \frac{p(\boldsymbol{x}^\star|1, \boldsymbol{\mu}_1, \hat{\psi}_1)\, \bar{\sigma}_1}{p(\boldsymbol{x}^\star|0, \boldsymbol{\mu}_0, \hat{\psi}_0)\, \bar{\sigma}_0 \;+\; p(\boldsymbol{x}^\star|1, \boldsymbol{\mu}_1, \hat{\psi}_1)\, \bar{\sigma}_1}$$

$$\propto \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathrm{d}\boldsymbol{\mu}_0 \mathrm{d}\boldsymbol{\mu}_1\, \mathrm{e}^{-\frac{1}{2\hat{s}_0^2}(\boldsymbol{\mu}_0 - (\frac{\hat{s}_0}{\hat{\psi}_0})^2 \sum_{i=1}^N \delta_{\sigma_i,0}\, \boldsymbol{x}_i)^2 - \frac{1}{2\hat{s}_1^2}(\boldsymbol{\mu}_1 - (\frac{\hat{s}_1}{\hat{\psi}_1})^2 \sum_{i=1}^N \delta_{\sigma_i,1}\, \boldsymbol{x}_i)^2} \Big[1 + \frac{p(\boldsymbol{x}^\star|0, \boldsymbol{\mu}_0, \hat{\psi}_0)\bar{\sigma}_0}{p(\boldsymbol{x}^\star|1, \boldsymbol{\mu}_1, \hat{\psi}_1)\bar{\sigma}_1}\Big]^{-1}$$

$$\text{with} \quad \hat{s}_0^{-2} = \hat{\alpha}_0 + \hat{\psi}_0^{-2} \sum_{i=1}^N \delta_{\sigma_i,0} \qquad \hat{s}_1^{-2} = \hat{\alpha}_1 + \hat{\psi}_1^{-2} \sum_{i=1}^N \delta_{\sigma_i,1}$$

Transforming the integration variables using

$$\boldsymbol{\mu}_0 = (\frac{\hat{s}_0}{\hat{\psi}_0})^2 \sum_{i=1}^N \delta_{\sigma_i,0}\, \boldsymbol{x}_i + \hat{s}_0 \boldsymbol{y}_0 \qquad \boldsymbol{\mu}_1 = (\frac{\hat{s}_1}{\hat{\psi}_1})^2 \sum_{i=1}^N \delta_{\sigma_i,1}\, \boldsymbol{x}_i + \hat{s}_1 \boldsymbol{y}_1$$

$$\text{where} \quad \boldsymbol{y}_0 \sim \mathcal{N}(0, \mathbb{I}_{d\times d}) \qquad \boldsymbol{y}_1 \sim \mathcal{N}(0, \mathbb{I}_{d\times d})$$

gives

$$p(1|\boldsymbol{x}^\star, D) = \frac{\bar{\sigma}_1 \hat{\psi}_0^d}{\bar{\sigma}_0 \hat{\psi}_1^d} \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \Big[\frac{\bar{\sigma}_1 \hat{\psi}_0^d}{\bar{\sigma}_0 \hat{\psi}_1^d} + \mathrm{e}^{\frac{1}{2}(\frac{\hat{s}_1}{\hat{\psi}_1})^2 \left(\frac{\boldsymbol{x}^\star}{\hat{s}_1} - \frac{\hat{s}_1}{\hat{\psi}_1^2} \sum_{i=1}^N \delta_{\sigma_i,1}\, \boldsymbol{x}_i - \boldsymbol{y}_1\right)^2 - \frac{1}{2}(\frac{\hat{s}_0}{\hat{\psi}_0})^2 \left(\frac{\boldsymbol{x}^\star}{\hat{s}_0} - \frac{\hat{s}_0}{\hat{\psi}_0^2} \sum_{i=1}^N \delta_{\sigma_i,0}\, \boldsymbol{x}_i - \boldsymbol{y}_0\right)^2} \Big]^{-1} \mathcal{D}\boldsymbol{y}_1\, \mathcal{D}\boldsymbol{y}_0$$

This is a $2d$-dimensional integral of the following form

$$p(1|\boldsymbol{x}^\star, D) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{1}{1 + \mathrm{e}^{a_1(\boldsymbol{y}_1 - \tilde{\boldsymbol{y}}_1)^2 - a_0(\boldsymbol{y}_0 - \tilde{\boldsymbol{y}}_0)^2 - d\log(C_{01}) + \log(E_{01})}}\, \mathcal{D}\boldsymbol{y}_1\, \mathcal{D}\boldsymbol{y}_0 \tag{2}$$

with

$$a_1 = \frac{1}{2}(\frac{\hat{s}_1}{\hat{\psi}_1})^2 \quad a_0 = \frac{1}{2}(\frac{\hat{s}_0}{\hat{\psi}_0})^2 \quad \tilde{\boldsymbol{y}}_1 = \frac{\boldsymbol{x}^\star}{\hat{s}_1} - \frac{\hat{s}_1}{\hat{\psi}_1^2} \sum_{i=1}^N \delta_{\sigma_i,1}\, \boldsymbol{x}_i \quad \tilde{\boldsymbol{y}}_0 = \frac{\boldsymbol{x}^\star}{\hat{s}_0} - \frac{\hat{s}_0}{\hat{\psi}_0^2} \sum_{i=1}^N \delta_{\sigma_i,0}\, \boldsymbol{x}_i \quad C_{01} = \frac{\hat{\psi}_0}{\hat{\psi}_1} \quad E_{01} = \frac{\bar{\sigma}_0}{\bar{\sigma}_1}$$

The MAP estimates for the hyperparameters ψ_σ and α_σ are found by minimising the log-likelihood in the usual way. We choose a flat prior for ψ_σ. The only information we have about α_σ is that it is the precision of the mean $\boldsymbol{\mu}_\sigma$, and ≥ 0. Therefore we can choose an entropic prior [39, 40, 41] for α_σ. This also has the advantage of rounding off the Bayesian hierarchy at hyperparameter level. The optimal hyperparameters used in (2) are:

$$\hat{\alpha}_\sigma = 0 \quad \hat{\psi}_\sigma^2 = \frac{N_\sigma \boldsymbol{\Sigma}_\sigma^2}{(N_\sigma - 1)d}$$

where

$$N_\sigma = \sum_{i=1}^N \delta_{\sigma_i,\sigma} \quad \langle \boldsymbol{x} \rangle_\sigma = \frac{1}{N_\sigma} \sum_{i=1}^N \delta_{\sigma_i,\sigma}\, \boldsymbol{x}_i \quad \langle \boldsymbol{x}^2 \rangle_\sigma = \frac{1}{N_\sigma} \sum_{i=1}^N \delta_{\sigma_i,\sigma}\, \boldsymbol{x}_i^2 \quad \boldsymbol{\Sigma}_\sigma^2 = \langle \boldsymbol{x}^2 \rangle_\sigma - \langle \boldsymbol{x} \rangle_\sigma^2$$

Evaluating integral $p(1|\mathbf{x}^\star, D)$ for finite d

We rotate the basis for the vectors $\boldsymbol{y}_1$ and $\boldsymbol{y}_0$ such that the first unit vector in the space of $\boldsymbol{y}_1$ (y_{11}) points towards $\tilde{\boldsymbol{y}}_1$, and the first unit vector in the space of $\boldsymbol{y}_0$ (y_{01}) points in the direction of $\tilde{\boldsymbol{y}}_0$. This results in a

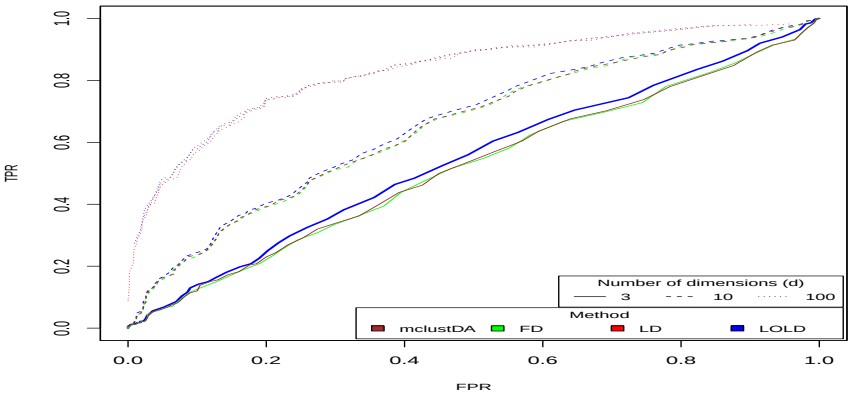

Figure 1: ROC curves for the simulated data sets using the different versions of our Bayesian method for $d = 3$ (solid line), $d = 10$ (dashed line), and $d = 100$ (dotted line).

formula in terms of the absolute values $\tilde{y}_1 = |\tilde{\boldsymbol{y}}_1|$, and $\tilde{y}_0 = |\tilde{\boldsymbol{y}}_0|$ alone:

$$p(1|\boldsymbol{x}^\star, D) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} W(Y_1)\, W(Y_0) \frac{1}{1+\mathrm{e}^{a_1 Y_1 - a_0 Y_0 + a_1(y_{11}-\tilde{y}_1)^2 - a_0(y_{01}-\tilde{y}_0)^2 - d\log(C) - \log(E)}} \mathrm{d}Y_1\, \mathrm{d}Y_0\, \mathcal{D}y_{11}\, \mathcal{D}y_{01} \tag{3}$$

where, for $d \geq 3$,

$$W(Y_\sigma) = \int_{-\infty}^{\infty} \mathcal{D}y_2 \ldots \mathcal{D}y_d\ \delta[Y_\sigma - \sum_{i=2}^{d} y_i^2] = \int_{-\infty}^{\infty} \frac{\mathrm{d}v}{2\pi} \frac{\mathrm{e}^{ivY_\sigma}}{(1+2iv)^{\frac{d-1}{2}}} = \frac{2^{-\frac{(d-1)}{2}}}{\Gamma(\frac{d-1}{2})}\, \theta(Y_\sigma)\ \ \mathrm{e}^{[-\frac{Y_\sigma}{2} + (\frac{d-3}{2})\log(Y_\sigma)]} \tag{4}$$

Substituting (4) into (3) gives

$$p(1|\boldsymbol{x}^\star, D) = \frac{2^{-(d-1)}}{\left(\Gamma(\frac{d-1}{2})\right)^2} \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{0}^{\infty}\int_{0}^{\infty} \frac{e^{-\frac{1}{2}(Y_1+Y_0) + (\frac{d-3}{2})\log(Y_1 Y_0)}}{1+\mathrm{e}^{a_1 Y_1 - a_0 Y_0 + a_1(y_{11}-\tilde{y}_1)^2 - a_0(y_{01}-\tilde{y}_0)^2 - d\log(C) - \log(E)}} \mathrm{d}Y_1\, \mathrm{d}Y_0\, \mathcal{D}y_{11}\, \mathcal{D}y_{01} \tag{5}$$

where (5) is 4-dimensional integral. Thus for $d \geq 3$ we have shown that solving the d-dimensional Bayesian integrals analytically has simplified our predictive distribution to a 4-dimensional integral (FD).

Evaluating integral $p(1|\mathbf{x}^\star, D)$ for large d

For large d, we can use the central limit theorem to write $Y_\sigma = (d-1) + Y_\sigma^\star\sqrt{2(d-1)}$, where $Y_\sigma^\star \sim \mathcal{N}(0,1)$. This gives

$$W(Y_\sigma) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathrm{d}Y_\sigma^\star\, \mathcal{D}Y_\sigma^\star\ \delta\Big[Y_\sigma - \Big((d-1) + Y_\sigma^\star\sqrt{2(d-1)}\Big)\Big] \tag{6}$$

Substituting (6) into (3) gives

$$p(1|\boldsymbol{x}^\star, D) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{1}{1+\mathrm{e}^{(a_1-a_0)(d-1) + \sqrt{2(d-1)}\sqrt{a_1^2+a_0^2}w}\ \mathrm{e}^{a_1(y_{11}-\tilde{y}_1)^2 - a_0(y_{01}-\tilde{y}_0)^2 - d\log(C) - \log(E)}} \mathcal{D}w\, \mathcal{D}y_{11}\, \mathcal{D}y_{01} \tag{7}$$

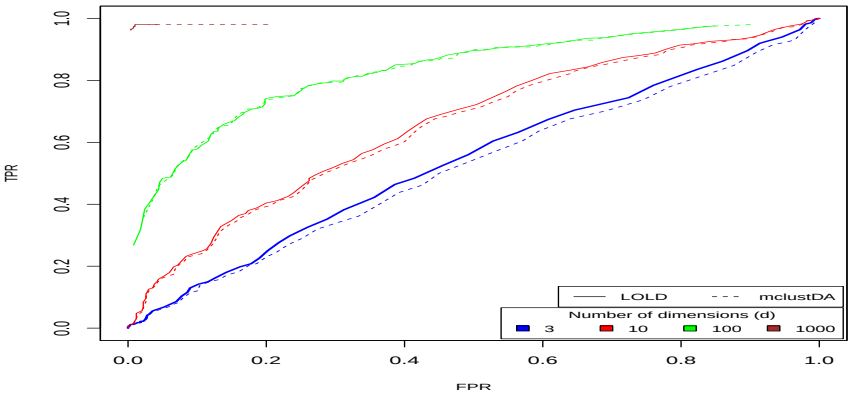

Figure 2: ROC curves for the simulated data sets $d = 3, 10, 100, 1000$, using LOLD (solid line) and *mclustDA* (dashed line).

where we have combined the sum $a_1 Y_1^\star - a_0 Y_0^\star$ into a single Gaussian variable w, with variance $a_1^2 + a_0^2$. Thus for large d we reduce the 4-dimensional integral to a 3-dimensional one (LD).

Evaluating integral $p(1|\mathbf{x}^\star, D)$ for leading order large d

If we only take the leading order terms (the other terms become negligible as d gets large), then (7) becomes a simple formula

$$p(1|\boldsymbol{x}^\star, D) = (1 + \mathrm{e}^{\,d\,\Lambda_{10}})^{-1}, \qquad \text{with} \quad \Lambda_{10} = a_1 - a_0 - \log(C) + \frac{a_1 \tilde{y}_1^2}{d} - \frac{a_0 \tilde{y}_0^2}{d}$$

Thus we have

$$p(1|\boldsymbol{x}^\star, D) = \begin{cases} 0 + \mathcal{O}(\mathrm{e}^{-d\,\Lambda_{10}}) & \text{if} \quad \Lambda_{10} > 0 \\ 1 - \mathcal{O}(\mathrm{e}^{-d\,|\Lambda_{10}|}) & \text{if} \quad \Lambda_{10} < 0 \end{cases}$$

which gives

$$p(1|\boldsymbol{x}^\star, D) = \theta(-\Lambda_{10}) \pm \mathcal{O}(\mathrm{e}^{-d\,|\Lambda_{10}|}) \tag{8}$$

Thus we have shown that for leading order large d the $2d$-dimensional predictive distribution simplifies to a simple formula (LOLD) in (8).

4 Simulation study

In this section, we apply our proposed Bayesian method, and *mclustDA* to simulated data sets. In order to compare our proposed Bayesian method to *mclustDA*, we simulated data sets consisting of $N = 100$ observations, and their corresponding classifications. The model complexity (the number of unknown parameters) was varied by increasing the dimension of each data set, from the lowest dimension that our method can be applied to, $d = 3$, to $d = 10, 100, 1000, 10000$. The last dimension cannot be handled by *mclustDA*, and is used to illustrate the power of our Bayesian method in very large dimensions that are typical of large biomedical data sets [11, 12, 28, 29, 30, 31, 32].

We simulated data sets whose classes had the same mean vector in order to test the predictive and classification power of our methods, since well-separated classes are easier to classify particularly in higher dimensions [18, 21]. The variances of the classes were set to be weakly different values, in order to allow one class to nest inside the space of the other. This allowed the classes to become more separated as the dimension increased, since the location of the peaks (mass) of the Gaussian distributions move away from each other whilst the width of each distribution increases much slower (until it becomes constant) as the dimension increases [21]. In contrast, where the mean vectors and variances are the same, the classes are indistinguishable [21].

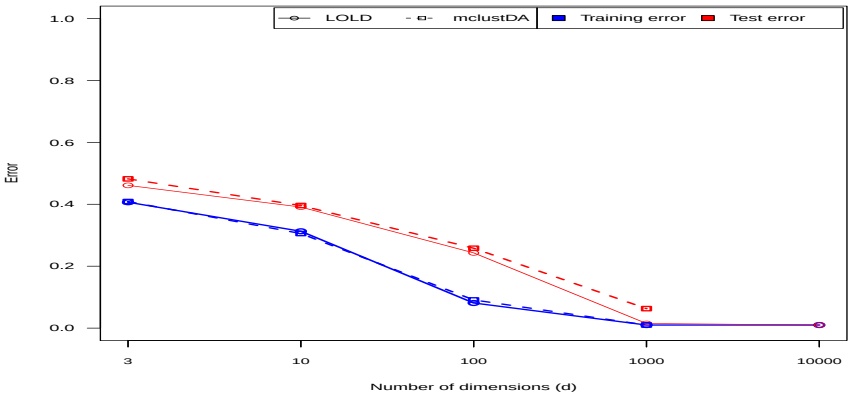

Figure 3: Error curves for the simulated data sets with $d = 3, 10, 100, 1000, 10000$ at a decision threshold of 0.5, using LOLD (solid line), and *mclustDA* (dashed line). The performance improves in high dimensions as the peaks (mass) of the Gaussian distributions move away from each other whilst the width of each distribution remains constant [21].

For each dimension, two classes, $\sigma = \{0, 1\}$, of observations were simulated according to $\boldsymbol{x}|\sigma \sim \mathcal{N}(\boldsymbol{\mu}_\sigma, \lambda_\sigma^2 \, \mathbb{I}_{d\times d})$, with $\mu_0 = \mu_1 = 0$, $\lambda_0 = 0.24$, and $\lambda_1 = 0.28$. The number of observations in each class was kept the same ($N_0 = N_1 = 50$). Classifications, $\sigma = \{0, 1\}$, were generated for the number of observations in each class respectively. The model parameters, class-imbalance parameters, and number of observations for each dimension are summarised in Table 1.

Table 1: Parameters of the simulated isotropic Gaussian data sets consisting of two distinct classes.

N	N_0	N_1	$\boldsymbol{\mu}_0$	$\boldsymbol{\mu}_1$	λ_0	λ_1	$\bar{\sigma}_0$	$\bar{\sigma}_1$
100	50	50	0	0	0.24	0.28	0.5	0.5

In order to determine when to use the different versions (FD, LD, LOLD) of our method, we applied all three versions to the simulated data sets with relatively low dimensions ($d = 3, 10$). Only LD and LOLD were applied to the simulated data sets with $d = 100$. *mclusDA* was applied as a reference. The results using a threshold of 0.5 are shown in Table 2. The receiver operating characteristic (ROC) curves [42] for all the versions and dimensions are displayed in Figure 1. We observe that LOLD preforms best which suggests that the leading order (constant) $\mathcal{O}(d)$ terms dominate the integrated terms which are $\mathcal{O}(1)$. Therefore we continue our analysis using LOLD.

The ROC curves for both methods in the different dimensions are shown in Figure 2. As expected, neither LOLD nor *mclustDA* perform well for $d = 3$ as they cannot distinguish between the classes. Both methods improve as the dimension of the data set increases given that it lives in subspaces with an effective dimensionality less than d [21, 18], although LOLD outperforms *mclustDA* on average in all dimensions. For $d = 10000$, LOLD performs with 99 per cent accuracy.

Using a classification threshold of 0.5, Figure 3 illustrates the average training and test error for both methods in different dimensions. We notice that for LOLD overfitting is generally reduced as the number of dimensions increases until it ceases for $d = 1000$. For *mclustDA*, overfitting is reduced at a slower rate but remains for $d = 1000$. Therefore we have demonstrated that on average, LOLD performs as well as or better than *mclustDA*, and reduces overfitting until it ceases for $d = 1000$.

Table 2: Average performance comparison for the different methods.

d	Method	Threshold	Accuracy	Computation time (mins)
3	*mclustDA*	0.5	0.52	0.01
	FD	0.5	0.52	32551.78
	LD	0.5	0.52	2214.55
	LOLD	0.5	0.54	0.00
10	*mclustDA*	0.5	0.60	0.01
	FD	0.5	0.61	57295.17
	LD	0.5	0.61	3897.88
	LOLD	0.5	0.61	0.00
100	*mclustDA*	0.5	0.74	0.09
	LD	0.5	0.76	4263.10
	LOLD	0.5	0.76	0.00
1000	*mclustDA*	0.5	0.94	9.14
	LOLD	0.5	0.99	0.00
10000	LOLD	0.5	0.99	0.07

5 Application to triple-negative breast cancer data

In this section, we apply our proposed Bayesian method, and *mclustDA* to a real example. The gene expression data set for triple-negative breast cancer patients (*exprset2*) from *de Rinaldis et al.* [29] has been used. The data set consists of $N = 176$ patients, who were treated at Guy's and St Thomas's Hospitals, London, UK between 1979 and 2001, and had at least 5.5 years follow-up. Each patient had $d = 22035$ gene expressions recorded. Patients with missing data or who were lost to follow up were excluded from our analysis. This yielded $N = 165$ patients to study.

Patients who survived for an interval of at least 5 years from initial diagnosis were designated as class $\sigma = 1$, the good prognosis group: GP(1). Patients who died from breast cancer within 5 years were designated as class $\sigma = 0$, the poor prognosis group: GP(0). The aim was to use the patients' gene expression data to predict which prognosis group they belong to. The total number of patients, the number of patients in each class, and label-imbalance parameters are summarised in Table 3.

Table 3: *exprset2* data set.

d	GP(0)	GP(1)	N	N_0	N_1	$\bar{\sigma}_0$	$\bar{\sigma}_1$
22035	$< 5yrs$	$\geq 5yrs$	165	41	124	0.25	0.75

In order to apply *mclustDA* to the *exprset2* data set, we reduced the number of dimensions a priori since *mclustDA* could not handle such large dimensions. Based on the hypothesis that genes with greater correlation with outcome are likely candidates for reporting prognosis [43, 44], the correlation coefficient of the expression for each gene, and outcome was calculated. The genes were then ranked-ordered according to the magnitude of the correlation coefficient [43, 44]. The highest ranked 100, 1000 and 10000 genes were chosen. *mclustDA* was applied to the first two reduced *exprset2* data sets consisting of $N = 165$ patients and $d = 100$ genes, and $N = 165$ patients and $d = 1000$ genes respectively. LOLD was applied to all three reduced data sets, and the full data set. The results using a decision threshold of 0.5 are shown in Table 4. LOLD performed as well as *mclustDA*, and was more efficient for $d = 100$. For $d = 1000$, LOLD was both more accurate and efficient

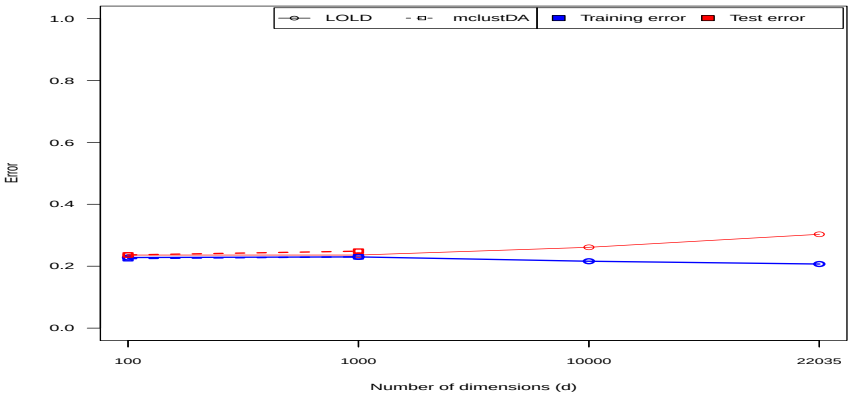

Figure 4: Error curves for the exprset2 data sets $d = 100, 1000, 10000, 22035$ at a decision threshold of 0.5, using LOLD (solid line), and *mclustDA* (dashed line).

than *mclustDA*. In higher dimensions, overfitting causes the predictive accuracy of LOLD to decrease. Using a classification threshold of 0.5, Figure 4 illustrates the training, and test error for both methods. Compared to *mclustDA*, LOLD reduces overfititng for the two lower dimensions. For the two higher dimensions, LOLD is clearly overfitting although the optimal predictive accuracy only decreases by 6 per cent.

Table 4: *exprset2* data set results.

Method	d	Threshold	Accuracy	Computation time (mins)
mclustDA	100	0.5	0.76	0.09
	1000	0.5	0.75	11.36
LOLD	100	0.5	0.76	0.00
	1000	0.5	0.76	0.03
	10000	0.5	0.74	0.40
	22035	0.5	0.70	1.67

Additionally, we now know a posteriori that the predictive information lies in the first 1000 rank-ordered genes for this data set. This would not have been possible with *mclustDA* as it was less accurate for these 1000 genes, and could not handle the first 10000 rank-ordered genes. This information may have been lost using a dimension reduction technique a priori, since it would require a user-defined arbitrary cut-off which may have been less than the first 1000 rank-ordered genes.

6 Discussion

We have proposed a Bayesian binary outcome prediction method that can be applied to any arbitrary number of dimensions without any approximations at parameter level. Our method overcomes the computational inability associated with other methods in large dimensions without sacrificing any information from the data set. Although three prediction formulas, (FD, LD, LOLD) have been developed, we have shown that it suffices to use the most efficient version, LOLD, for any arbitrary dimension, without compromising predictive performance. We compared our method to *mclusDA* [2, 17] due to its Bayesian approach, and wide application [3, 4, 5, 6, 7]. For data sets where the methods produce the same results, our method is more efficient. For high-dimensional data sets, our Bayesian method is superior in both accuracy and efficiency. Our method also succeeds in reducing overfitting although it remained in very high-dimensions.

Our method can be seen as an extension for *mclustDA*, in a homogeneous binary class setting, where we take the uncertainty at parameter level into account instead of estimating the Bayesian integrals using EM. However the methods differ in their application to the training set. We use a supervised method by applying our Bayesian method to training set using both the training data and outcome information whilst *mclustDA* applies an unsupervised model-based clustering algorithm to the training data only. A limitation of our approach is that we must set the number of classes in the data set, and select the optimal model a priori, whilst *mclustDA* calculates this from the data. We also limited our approach to only two homogeneous classes. A natural extension would be to expand our approach to consider more than two classes. Our method could also extend to heterogeneous classes. The would entail fitting a more complex GMM to each class. We would also need to consider Bayesian model selection in order to choose the optimal parametrisation for each GMM. The uncertainty in model selection, and at hyperparameter level could also be incorporated into our frame work. However the risk of heavily overfitting remains for all these ideas, as the number of unknown parameters, and hyperparameters will increase considerably.

Finally, we have only compared our method to the *mclustDA*. Many other model-based discriminant analysis methods exist that overcome computational inability in high-dimension [15, 20]. These include using variable selection [45, 6], and combining subspace clustering with constrained and parsimonious models [14, 20]. However, these methods may still suffer from loss of information, and/or underestimated uncertainty at parameter level. Overfitting remains where the imbalance between the number of observations and dimensions remains high, even after reducing the dimensionality of the data set. The method we propose overcomes these problems by reducing the number of dimensions analytically, whilst also taking the uncertainty at parameter level into account. No information is lost during this process, and we postpone the use of computer resources until the resulting formulae can no longer be simplified.

References

[1] Hastie T, Tibshirani R. Discriminant analysis by gaussian mixtures. *Journal of the Royal Statistical Society, Series B* 1996; **58**:155–176.

[2] Fraley C, Raftery AE. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* 2002; **97**(458):611–631, doi:10.1198/016214502760047131. URL http://www.tandfonline.com/doi/abs/10.1198/016214502760047131.

[3] Dean N, Murphy TB, Downey G. Using unlabelled data to update classification rules with applications in food authenticity studies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2006; **55**(1):1–14, doi:10.1111/j.1467-9876.2005.00526.x. URL http://dx.doi.org/10.1111/j.1467-9876.2005.00526.x.

[4] Fraley C, Raftery AE. Model-based methods of classification: Using the mclust software in chemometrics. *JOURNAL OF STATISTICAL SOFTWARE* JAN 2007; **18**(6).

[5] Iverson AA, Gillett C, Cane P, Santini CD, Vess TM, Kam-Morgan L, Wang A, Eisenberg M, Rowland CM, Hessling JJ, *et al.*. A single-tube quantitative assay for mrna levels of hormonal and growth factor receptors in breast cancer specimens. *The Journal of Molecular Diagnostics* 2009; **11**(2):117 – 130, doi:http://dx.doi.org/10.2353/jmoldx.2009.080070. URL http://www.sciencedirect.com/science/article/pii/S1525157810602176.

[6] Murphy TB, Dean N, Raftery AE. Variable selection and updating in model-based discriminant analysis for high dimensional data with food authenticity applications. *The annals of applied statistics* 2010; **4**(1):396.

[7] Andrews J, McNicholas P. Model-based clustering, classification, and discriminant analysis via mixtures of multivariate t-distributions. *Statistics and Computing* 2012; **22**(5):1021–1029, doi:10.1007/s11222-011-9272-x. URL http://dx.doi.org/10.1007/s11222-011-9272-x.

[8] Lachenbruch PA, Mickey MR. Estimation of error rates in discriminant analysis. *Technometrics* 1968; **10**(1):pp. 1–11. URL http://www.jstor.org/stable/1266219.

[9] Stone M. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society. Series B (Methodological)* 1974; **36**(2):pp. 111–147. URL http://www.jstor.org/stable/2984809.

[10] Meek C, Thiesson B, Heckerman D, Kaelbling P. The learning-curve sampling method applied to model-based clustering. *Also in AI and Statistics*, 2001; 397–418.

[11] Clarke R, Ressom HW, Wang A, Xuan J, Liu MC, Gehan EA, Wang Y. The properties of high-dimensional data spaces: implications for exploring gene and protein expression data. *Nat Rev Cancer* Jan 2008; **8**(1):37–49. URL http://dx.doi.org/10.1038/nrc2294.

[12] Michiels S, Kramar A, Koscielny S. Multidimensionality of microarrays: Statistical challenges and (im)possible solutions. *Mol Oncol* Apr 2011; **5**(2):190–196. URL http://linkinghub.elsevier.com/retrieve/pii/S1574789111000184?showall=true.

[13] Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B* 1977; **39**(1):1–38.

[14] Bouveyron C, Girard S, Schmid C. High-dimensional discriminant analysis. *Communications in Statistics - Theory and Methods* 2007; **36**(14):2607–2623, doi:10.1080/03610920701271095. URL http://www.tandfonline.com/doi/abs/10.1080/03610920701271095.

[15] McNicholas PD. On model-based clustering, classification, and discriminant analysis. *Journal of the Iranian Statistical Society* 2011; **10(2)**:181–199.

[16] Jolliffe IT. *Principal Component Analysis (2nd edn)*. Springer, 2002.

[17] Fraley C, Raftery AE, Murphy TB, Scrucca L. mclust version 4 for r: Normal mixture modeling for model-based clustering, classification, and density estimation 2012.

[18] Scott DW, Thompson JR. Probability density estimation in higher dimensions. *Computer Science and Statistics: Proceedings of the Fifteenth Symposium on the Interface*, vol. 528, North-Holland, Amsterdam, 1983; 173–179.

[19] Chang WC. On using principal components before separating a mixture of two multivariate normal distributions. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 1983; **32**(3):pp. 267–275. URL http://www.jstor.org/stable/2347949.

[20] Bouveyron C, Brunet-Saumard C. Model-based clustering of high-dimensional data: A review. *Computational Statistics & Data Analysis* 2014; **71**(0):52 – 78, doi:http://dx.doi.org/10.1016/j.csda.2012.12.008. URL http://www.sciencedirect.com/science/article/pii/S0167947312004422.

[21] Bishop CM. *Pattern Recognition and Machine Learning*. Springer, 2006.

[22] Berge L, Bouveyron C, Girard S. Hdclassif: An r package for model-based clustering and discriminant analysis of high-dimensional data. *Journal of Statistical Software* 1 2012; **46**(6):1–29. URL http://www.jstatsoft.org/v46/i06.

[23] Bellman R. *Dynamic Programming*. Princeton University Press, 1957.

[24] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria 2013. URL http://www.R-project.org/.

[25] Ripley BD. *Pattern Recognition and Neural Networks*. Cambridge University Press, 1996.

[26] Duda R, Hart P, Stork D. *Pattern Classification (2nd edn)*. Wiley, 2001.

[27] McLachlan G, Peel D, Bean R. Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics & Data Analysis* 2003; **41**(34):379 – 388, doi:http://dx.doi.org/10.1016/S0167-9473(02)00183-4. URL http://www.sciencedirect.com/science/article/pii/S0167947302001834, recent Developments in Mixture Model.

[28] Wang Y, Miller DJ, Clarke R. Approaches to working in high-dimensional data spaces: gene expression microarrays. *Br J Cancer* Feb 2008; **98**(6):1023–1028. URL http://dx.doi.org/10.1038/sj.bjc.6604207.

[29] de Rinaldis E, Gazinska P, Mera A, Modrusan Z, Fedorowicz G, Burford B, Gillett C, Marra P, Grigoriadis A, Dornan D, *et al.*. Integrated genomic analysis of triple-negative breast cancers reveals novel microrna s associated with clinical and molecular phenotypes and sheds light on the pathways they control. *BMC Genomics* 2013; **14**(1):643, doi:10.1186/1471-2164-14-643. URL http://www.biomedcentral.com/1471-2164/14/643.

[30] Network TCGAR. Comprehensive molecular portraits of human breast tumours. *Nature* Oct 2012; **490**(7418):61–70, doi:10.1038/nature11412.

[31] Network TCGAR. Integrated genomic analyses of ovarian carcinoma. *Nature* 06 2011; **474**(7353):609–615. URL http://dx.doi.org/10.1038/nature10166.

[32] Barretina J, Caponigro G, Stransky N, Venkatesan K, Margolin AA, Kim S, Wilson CJ, Lehar J, Kryukov GV, Sonkin D, *et al.*. The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature* 03 2012; **483**(7391):603–307. URL http://dx.doi.org/10.1038/nature11003.

[33] Dasgupta A, Raftery AE. Detecting features in spatial point processes with clutter via model-based clustering. *Journal of the American Statistical Association* 1998; **93**(441):pp. 294–302. URL http://www.jstor.org/stable/2669625.

[34] Fraley C, Raftery AE. How many clusters? which clustering method? answers via model-based cluster analysis. *The Computer Journal* 1998; **41**:578–588.

[35] Banfield JD, Raftery AE. Model-based Gaussian and non-Gaussian clustering. *Biometrics* 1993; **49**(3):803–821.

[36] Schwarz G. Estimating the Dimension of a Model. *The Annals of Statistics* 1978; **6**(2):461–464, doi:10.2307/2958889. URL http://dx.doi.org/10.2307/2958889.

[37] Wehrens R, Buydens LM, Fraley C, Raftery AE. Model-based clustering for image segmentation and large datasets via sampling. *Journal of Classification* 2004; **21**(2):231–253, doi:10.1007/s00357-004-0018-8. URL http://dx.doi.org/10.1007/s00357-004-0018-8.

[38] Fraley C, Raftery A, Wehrens R. Incremental model-based clustering for large datasets with small clusters. *Journal of Computational and Graphical Statistics* 2005; **14**(3):529–546, doi:10.1198/106186005X59603. URL http://amstat.tandfonline.com/doi/abs/10.1198/106186005X59603.

[39] Rodriguez CC. Entropic priors for discrete probabilistic networks and for mixtures of gaussians models. *arXiv preprint physics/0201016* 2002; .

[40] Caticha A, Preuss R. Maximum entropy and bayesian data analysis: Entropic prior distributions. *Physical Review E* 2004; **70**(4):046 127.

[41] Neumann T. Bayesian inference featuring entropic priors. *27th Int. Work. on Bayesian Inf. and Max. Ent* 2007; **954**(1):283–292.

[42] Gerds TA, Cai T, Schumacher M. The performance of risk prediction models. *Biometrical Journal* 2008; **50**(4):457–479, doi:10.1002/bimj.200810443. URL http://dx.doi.org/10.1002/bimj.200810443.

[43] van 't Veer LJ, Dai H, van de Vijver MJ, He YD, Hart AAM, Mao M, Peterse HL, van der Kooy K, Marton MJ, Witteveen AT, *et al.*. Gene expression profiling predicts clinical outcome of breast cancer. *Nature* Jan 2002; **415**(6871):530–536. URL http://dx.doi.org/10.1038/415530a.

[44] van de Vijver MJ, He YD, van 't Veer LJ, Dai H, Hart AA, Voskuil DW, Schreiber GJ, Peterse JL, Roberts C, Marton MJ, *et al.*. A gene-expression signature as a predictor of survival in breast cancer. *New England Journal of Medicine* 2002; **347**(25):1999–2009, doi:10.1056/NEJMoa021967. URL http://www.nejm.org/doi/full/10.1056/NEJMoa021967, pMID: 12490681.

[45] Raftery AE, Dean N. Variable selection for model-based clustering. *Journal of the American Statistical Association* 2006; **101**(473):168–178, doi:10.1198/016214506000000113. URL http://www.tandfonline.com/doi/abs/10.1198/016214506000000113.